BOOK REVIEWS

Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg, eds., ***Scientific Collaboration and Collective Knowledge: New Essays.*** New York: Oxford University Press (2017), 240 pp., $85.00 (cloth).

How can and should the social dimensions of science contribute to science's ability to acquire knowledge? Given the central role that collaborative practices, sharing and publishing practices, rewards and incentives, disagreement and consensus formation play in science, questions about how these social dimensions positively or negatively affect scientific knowledge generation are central to the philosophy of science. Yet surprisingly such questions have been little studied in philosophy of science, where the focus has been almost exclusively on questions about the nature of scientific theories and their relationship to evidence, abstracting away from scientists and their social interactions.

The volume *Scientific Collaboration and Collective Knowledge*, edited by Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg, aims to remedy this. Containing nine essays on epistemological issues related to scientific collaboration, it is among the first collections of work in the social epistemology of science and forms part of a trend of increasing interest in this type of questions.

Perhaps the strongest aspect of this volume is that it brings together authors with very different methodological approaches, who at the same time share an interest in normative questions about the social dimensions of science. That a diversity of methodological approaches is epistemically valuable has long been a central claim of the social epistemology of science (e.g., Paul Feyerabend, *Against Method* [London: New Left Books, 1975]), so it is nice to see this thesis self-exemplified in the field and this volume in particular.

The volume brings together social epistemologists whose methodological approach relies on formal modeling (using analytical methods or simulations to obtain results from such models) and those who rely on detailed case studies. Authors taking a formal approach tend to focus on general claims, intended to apply across different fields of science (including the social sciences and the humanities), whereas analysis of case studies tends to bring out unique features of each field and each case. The two approaches—one emphasizing commonalities and the other emphasizing differences—are complementary, and both benefit from being brought together.

The nine chapters of the volume may be sorted into four groups based on the type of question under consideration: Why do scientists choose to collaborate?

How and when should scientists share their findings? How has the fact of increased collaboration affected peer review and publication practices? And what may a scientific community be said to (collectively) believe about a question or a set of questions? The first question—What motivates scientists to collaborate and what do they get out of it?—is addressed by Kevin Zollman's, Ryan Muldoon's, and Justin Bruner and Cailin O'Connor's chapters.

Zollman's contribution builds on a model due to Venkatesh Bala and Sanjeev Goyal ("A Noncooperative Model of Network Formation," *Econometrica* 68, no. 5 [2000]: 1181–1229) in which scientists form an epistemic network by individually constructing links that provide them with valuable information. Zollman refers to this as learning to collaborate, but he has in mind a relatively impoverished sense of collaboration in which information flows in one direction.

Most of the chapter is devoted to evaluating the performance of scientists when they create their network using a very simple learning rule known as probe and adjust. The community turns out to perform best when scientists consider changing their strategy relatively infrequently. Lowering the cost of collaboration and increasing the size of the community turn out to be harmful as measured by the proportion of time the community spends in the socially optimal state, but these results reverse when the metric of interest is the average performance of the community, which is arguably more important in practice.

Muldoon sets out to answer the following questions: Why do scientists choose to collaborate? Why is specialization on the rise? And why do certain disciplines seem to take over others? He argues that existing models in the social epistemology of science should be inverted to answer these questions. Rather than thinking of scientists as having particular research problems and going in search of the tools, skills, or methods to solve them, Muldoon suggests we think of scientists as having invested in certain skills and seeking problems to apply them to.

A rational scientist aims to get as much out of these investments as possible. The prevalence of collaboration is explained, according to Muldoon, by the increasing rarity of problems that can be solved by a single scientist or lab. Specialization is explained by the growth of science and the fact that science works most efficiently when individuals differentiate their skill sets. Scientists will take over other fields when their skills are relatively transferable and relatively underpaid in their original field.

Since Muldoon provides neither a formal model nor detailed case studies, his explanations remain at a fairly broad-brush level. Future work could flesh out the details of Muldoon's arguments as well as reveal whether this way of thinking about the social epistemology of science is a fruitful source of new hypotheses or predictions.

The chapter by Bruner and O'Connor uses evolutionary game theory to understand how scientific communities can develop discriminatory collaborative norms. They model the division of credit among collaborating scientists as a bargaining game, with the scientists learning to adopt successful strategies in accordance with the replicator dynamics.

They make two key claims: first, that when the population consists of groups of different sizes, collaborative norms are likely to develop that discriminate against minority groups, even though no explicit preference for or against discrimination is built into the model, and second, that among equal-sized groups discriminatory norms may also evolve when there are differences in bargaining power between the groups. Such differences probably exist between senior and junior academics. Insofar as senior academics face a lower cost when a collaborative project fails (this is cashed out in three different ways), they are likely to develop collaborative norms that discriminate against junior academics. This chapter represents formal social epistemology at its best, as the authors advance an original hypothesis on a socially relevant issue using formal results that are convincingly shown to be robust.

The second set of contributions asks what and when scientists should share with each other. There exists a social norm in science that calls on scientists to share their work widely, known as the communist norm. Michael Strevens's chapter asks when it is rational to conform to this norm for a scientist interested in maximizing the credit that accrues to her in virtue of her discoveries.

Strevens's "Hobbesian vindication" claims that scientists should be willing to sign a binding contract that commits them to sharing all their findings. Although he does not put it this way, Strevens effectively claims that in sharing their results scientists face a Prisoners Dilemma: it would be best if everyone shared, but each individual has an incentive not to share.

While I am on record disagreeing with Strevens's claim that individual scientists do not have an incentive to share (Remco Heesen, "Communism and the Incentive to Share in Science," *Philosophy of Science* 84, no. 4 [2017]: 698–716), his argument that it is best for everyone if everyone shares is independent of that. In particular, his defense of the "Marxian precept" that fair sharing requires more from high-powered research programs is interesting and deserves further uptake.

Staffan Angere and Erik J. Olsson are interested in the question whether communication is good for science, especially in light of (other) work by Venkatesh Bala and Sanjeev Goyal ("Learning from Neighbours," *Review of Economic Studies* 65, no. 3 [1998]: 595–621) and Kevin J. S. Zollman ("The Epistemic Benefit of Transient Diversity," *Erkenntnis* 72, no. 1 [2010]: 17–35) suggesting that denser communication networks are not always better. Angere and Olsson study this question in a model that is considerably different from Bala and Goyal's and Zollman's, especially in that agents are quite sophisticated in updating their beliefs through Bayesian conditionalization.

A key assumption made by Angere and Olsson is source independence, which means that agents treat communications from different agents as probabilistically independent. They also treat multiple communications from the same agent as probabilistically independent. This leads to a potential problem of double counting evidence, as one or a few signals from the world may be communicated repeatedly by one or more agents.

Angere and Olsson formulate conditions to reduce double counting. Whether more communication is epistemically beneficial in their model turns out to depend on these conditions. However, the informal gloss that scientists should exercise some restraint in publishing seems to capture the commonalities between the different versions of the model studied.

This conclusion—"publish late, publish rarely," as the authors put it—seems to agree with Zollman's "Epistemic Benefit of Transient Diversity." Whether and how this conclusion can be taken from an interesting suggestion to an actionable policy proposal remains an open question (Sarita Rosenstock, Justin Bruner, and Cailin O'Connor, "In Epistemic Networks, Is Less Really More? *Philosophy of Science* 84, no. 2 [2017]: 234–52; Daniel Frey and Dunja Šešelja, "What Is the Epistemic Function of Highly Idealized Agent-Based Models of Scientific Inquiry?" *Philosophy of the Social Sciences* 48, no. 4 [2018]: 407–33).

The next two contributions explore the consequences of increased collaboration, focusing especially on questions about responsibility for collaborative work. In their chapter, Bryce Huebner, Rebecca Kukla, and Eric Winsberg describe "radically collaborative research." In contrast to other forms of scientific collaboration, they argue, radically collaborative research projects are essentially epistemically distributed. They involve a great number of contributors making epistemic decisions (e.g., about how to apply or adapt a research protocol or how to trade off false positives and false negatives) that are impossible to coordinate in advance and that will not necessarily add up coherently.

As a result, such work has no author in the traditional sense: one or more people who, individually, can take full responsibility for the validity, integrity, and coherence of the finished product. The nature and scale of the problems raised when no individual(s) can be held accountable for a particular piece of collaborative research are left implicit by Huebner, Kukla, and Winsberg but are brought out by K. Brad Wray in the next chapter. Huebner, Kukla, and Winsberg outline a potential solution in a different article (Eric Winsberg, Bryce Huebner, and Rebecca Kukla, "Accountability and Values in Radically Collaborative Research," *Studies in History and Philosophy of Science* A 46

[2014]: 16–23). This requires one or more authors to explicitly form a theory of the (social) processes involved in the research and take responsibility for their epistemic consequences.

The first half of Wray's contribution summarizes his views on the role of collective belief in science. He argues that views expressed by collaborating scientists should be construed as collective acceptance, where acceptance differs from belief in that it is voluntary and may be pragmatically motivated.

In the second half, Wray considers the question of accountability also raised by Huebner, Kukla, and Winsberg. He points out that if no one is individually accountable for the results of collaborative scientific research, in cases of error or misconduct blame falls only on the collective. This significantly reduces the scope for real-world consequences. Journals and their editors have yet to find a satisfactory way to address this problem.

The increasing size and scope of collaborative research also raises problems for refereeing practices, of which Wray outlines two. On the one hand collaborations can grow so large that it is hard to find referees who are not authors on a project. On the other hand large collaborations may create situations in which no single scientist has the combination of expertise to referee every aspect of a project. Wray tentatively suggests collaborating teams of referees as a solution. While the additional time and cost associated with such a proposal may prove prohibitive, Wray's practice-oriented approach and his openness to innovation in peer review is to be welcomed.

The final two chapters are focused exclusively on the issue of collective belief also touched on by Wray. Can a group (e.g., a scientific community) be said to have beliefs? If so, under what circumstances? And when these circumstances apply, what is the content of the collective belief?

Denis Bonnay considers how to attribute collective beliefs to unorganized groups, where there is no notion of group members actively accepting or adopting a collective belief. Judgment aggregation theory has shown that attributing a single set of collective beliefs to a diverse group quickly leads to paradoxical outcomes. Bonnay's suggestion is to identify clusters within larger groups. According to Bonnay's motivating example, French people's beliefs about transgenic seeds are too diverse for a single set of collective beliefs but can plausibly be divided into a small number of clusters, for example, "Deep Greens," "Scientists," and "Skeptics," such that there is broad agreement within each cluster.

The rest of the article outlines a formal framework for studying "clustering methods," which identify clusters and apply judgment aggregation within each cluster. Assuming the clustering step satisfies some basic requirements, Bonnay shows that the conditions anonymity, systematicity, and unanimity preservation jointly characterize majoritarian aggregation (or equivalently, minimizing Hamming distance) for the judgment aggregation step.

The implicit point is that clustering should be used to identify subgroups whose beliefs are similar enough that majoritarian aggregation gives reasonable results. There is a nice convergence here with recent work by Liam Kofi Bright, Haixin Dang, and Remco Heesen ("A Role for Judgment Aggregation in Coauthoring Scientific Papers," *Erkenntnis* 83, no. 2 [2018]: 231–52), who argue that scientific collaborations should be considered successful only if their published work can be represented as the result of majoritarian aggregation.

In the final chapter, Carlo Martini and Jan Sprenger review the literature on opinion aggregation. The central question they aim to answer is under what circumstances an egalitarian model—one in which individuals' opinions are treated as interchangeable—should be preferred to one that aims to reflect differences in expertise among individuals. Guided by this question, the authors restrict their attention to cases in which a group estimate of a single probability or quantity is needed (setting aside cases in which a binary proposition needs to be judged or multiple related propositions are at issue).

A prominent approach to such cases relies on taking weighted averages of individuals' opinions. Martini and Sprenger endorse this as a reasonable (but not uncontroversial) solution. But what weights should be assigned to individuals' opinions? The central question here takes the form of whether each individual should receive equal weight or whether individuals should receive differential weights based on expertise.

From an epistemic perspective, this depends on how confident we are in our judgments of expertise. The authors report a number of results to the effect that weighted averaging outperforms egalitarian averaging under relatively modest conditions. For example, a sufficient condition is for each expert's weight to be between that expert's optimal weight and $1/n$. They conclude that differential weighting can be epistemically beneficial when the right kind of track record can be established for the experts.

Summing up, this volume provides a broad overview of the social epistemology of science, covering a wide range of subjects and methodologies. While the authors tend to focus on the philosophical conclusions of their work, it is clear that there are practical upshots as well. Angere and Olsson draw some tentative conclusions about what and when to publish, Wray considers a proposal for reforming peer review, and Klein and Sprenger discuss the circumstances under which differential weighting of opinions is epistemically beneficial, to name just three of the more explicit examples.

Exactly because they are social, the social dimensions of science are particularly suitable to intervention. Part of the goal of the social epistemology of science is to evaluate policy proposals or other interventions on their epistemic merits, that is, on the extent to which they will help or hinder science. While the conclusions of any one model or case study should not be considered decisive, the community of social epistemologists is ideally placed to make such normative judgments.

The volume is one of very few collections of work in this field. As such it has a potential role to play in introductory courses on the social epistemology of science, although it would have to be supplemented. In providing a state-of-the-art work, it will be of interest to anyone with research interests in the social epistemology of science. Even for those without such interests, it provides a model of how different methodologies can be fruitfully combined toward shared philosophical goals.

Remco Heesen, University of Western Australia

Carl Gillett, ***Reduction and Emergence in Science and Philosophy.*** Cambridge: Cambridge University Press (2016), 400 pp., $74.90 (cloth), $34.90 (paper).

Carl Gillett is convinced that current discussions about reduction and emergence, especially in philosophy, remain mired in an early twentieth-century understanding of what is at stake. That dispute was about the scope of compositional explanations across the sciences. Except for residual worries about, for example, phenomenal consciousness, the sciences throughout the twentieth century settled that dispute. “Everyday” reductionism won; all the entities, properties, relations, and processes in the world either are fully composed entities or are the most basic constituents that make things up. A new reductionism-emergentism dispute, however, is now raging in the sciences. It is over the nature of composition itself and the determinative relations between (some) composed entities and the powers of their parts. Philosophers have mostly missed this new dispute. Their current consensus, that “reductionism” is dead and “nonreductive physicalism” triumphant, results from two errors. The first is a widespread but mistaken acceptance of a “semantic” account of reduction(ism), rooted in logical positivism but still holding sway. The second is a “metaphysics for science” practice, as opposed to a “metaphysics of science” one more common among “working scientists” (Gillett’s term). By following Gillett in adopting the latter and clearly articulating the complex nature of “compositional explanation” at work throughout the sciences, we learn that there are currently three “live” positions within the contemporary scientific reduction-emergence dispute, two reductionist and one emergentist. As yet no one of these positions can claim a single successful case from science that fully meets Gillett’s explicit “threshold” for empirical adequacy (although his sympathies are with scientific emergentists).